

**Mass Storage at NSA****Michael F. Shields**

Department of Defense (J111)  
9800 Savage Road  
Ft. Meade, Maryland 20755-6000

Phone: 301-688-5065

Fax: 301-688-6152

**1.0 Introduction**

The need to manage large amounts of data on robotically controlled devices has been critical to the mission of this Agency for many years. In many respects this Agency has helped pioneer with their industry counterparts, the development of a number of products long before these systems became commercially available. Numerous attempts have been made to field both robotically controlled tape and optical disk technology and systems to satisfy our tertiary storage needs. Custom developed products were architected, designed and developed with our vendor partners over the past two decades to field workable systems to handle our ever increasing storage requirements. Many of the attendees of this symposium are familiar with some of the older products, such as: the Braegen Automated Tape Libraries (ATLs), the IBM 3850, the Ampex TeraStore, just to name a few. In addition, we embarked on an in-house development of a shared disk input/output support processor to manage our every increasing tape storage needs. For all intents and purposes, this system was a file server by current definitions which used CDC Cyber computers as the control processors. It served us well and was just recently removed from production usage.

**2.0 Project ABUNDANT**


In 1986, our existing storage systems were reaching the end of their useful lives. The STK Automated Library System (ACL) had not yet been developed and there were no existing Commercial Off the Shelf (COTS) products available from the vendor community that could satisfy our capacity and performance growth requirements. The Agency had an evolving requirement to field a 1015 bit robotic tape system that would interface with our existing and planned client supercomputer and server systems over our existing and planned networks. In addition the system had to occupy as small a physical space as possible. This system was to be a file server, use COTS components to the maximum extent possible, and be extensible by using a modular architecture.

Our first step was to fund known mass storage vendors to conduct independent concept definition studies to see if such a system could be built and more importantly would it be possible to achieve all of our program goals. Five vendors participated in this evolution and it was the unanimous belief that such a system could be built for an affordable price. A rather comprehensive set of Performance Requirement Specifications (PRS) were developed and an RFP was created. A competitive acquisition was held and as a result Project ABUNDANT was brought to life; in September of 1988, a contract was awarded and developmental work started on the Mass Storage Library (MSL).

A full time Program Management Office (PMO) was established to manage this multi-year effort. Several changes occurred over the course of this effort, resulting from both internal and external forces. The most important internal changes resulted from changing requirements and from a recognition that commercial products were starting to emerge from the vendor community. As a result, the program was restructured to take advantage of these factors, and the contract was changed to a two phased delivery of production equipment. In the summer of

# ABUNDANT PROGRAM

## R1 MSL ARCHITECTURE



1991, the first instantiation was installed and tested. It was based upon the architecture depicted in Figure One and employed a Convex control processor, the STK silo, and custom file server software developed by our integration contractor. The purpose of the interim system was to test the proof of concept and functionality of the product, and most importantly, to develop lessons learned which would help shape improvements for the larger 1015 bit system. The system was tested with Cray client systems during late 91 and early 92 and the desired lessons learned were captured.

The second release would also use a Convex control processor, but would have more functionality and increased performance. To handle the marked increase in capacity, an aisle based robotic tape archive with the desired modularity and capacity was developed. Built under a subcontract, this product was designed, developed, and tested by our integration contractor. This archive would use the ER90 D2 helical scan recorders and would be fully compliant with all of our stated goals and requirements. As we all know, what started as a custom development has now become a commercial product known as E-MASS. In addition, varying size robotic libraries are now commercially supported which include the STK Silo, the Odetics Data Tower (6 TBs) as well as our Data Library (150 TBs).

Taking advantage of the rapid expansion of commercially available Mass Storage product offerings, the Agency began restructuring the ABUNDANT program. In late FY92 a separate contract was awarded to field a network attached Volume Server based on the D2 recorders and the Odetics Data Tower. This system would use HiPPI switch technology to support tertiary tape storage needs of the recently acquired Cray YMP C-90 computer system used by our research organization. Once again it should be noted that we were acquiring a COTS, vice custom system. In addition, it was determined that the File Server EMASS product was mature enough so that the current ABUNDANT contract type (for the second release) could be changed from developmental to firm fixed price. Other operational changes have allowed this system to be utilized as a shared resource vice a dedicated system to a specific user group. Finally, to accelerate the fielding of the file server system, we decided to first implement it as a Data Tower this summer, and to then field the larger Data Library in late FY94. These changes have now been contractually implemented and final planning is being conducted to provide for a smooth installation this fiscal year.


Figure Two contains a depiction of both the interim and final file server configurations. It is important to note that this architecture is totally modular, offers significant flexibility for future change and upgrades, and clearly satisfies our COTS, footprint, and flexibility program goals. Figure Three contains a similar depiction of the HiPPI network attached Volume Server Data Tower system. Our principal activities over the next year will be to perform significant testing, of both approaches to Mass Storage so as to determine their optimal employment. Detailed plans are being developed with our customer, support, and operations organizations to fully evaluate both products.


### **3.0 Current Environment**

I alluded to the appearance in the commercial marketplace of a wide range of products that have, in most instances, found their way into our current operational environment, since the inception of ABUNDANT. Let me outline some of these that are used in our daily computer production. Perhaps the STK Silo could best be described as today's Mass Storage System of choice at the high end of the spectrum. Numerous silos are employed for the Cray, Convex, IBM and Amdahl, and other high end processors that we utilize. Silos are used as volume servers and are usually clustered in groups of two or three. They are cross connected to insure high availability and permit data interchange.

Next, we use both the Metrum RSS 48 and 600 SVHS robotic tape systems for mid-range processors. All of these are used as file servers and run the AMASS commercial file server product. Other AMASS uses employ robotic 8mm tape and optical disk libraries. In addition, Exabyte robotic controlled 8mm systems (EXB-120s, 10i's, and carousels) are used as volume

# CONVEX EMASS1200 FileServ Interim - 1300 FileServ Final at IOC





servers principally to perform backup function. A few user groups employ the Epoch file server software to manage their files.

#### **4.0 Near Term Environment**

Later this summer, the first UniTree evaluation will occur. This test will use the Amdahl as the control processor and the STK silo as the robotic library. Another user group is acquiring the TriPlex STX controller and the Sony ID-1 robotic library for use with a Cray YMP system. As previously stated, both the EMASS Data Tower FileServ and VolServ systems will be installed this year for evaluation. The Data Library EMASS FileServ product will replace the Data Tower in 1994. In addition several IBM 9570 RAID disk devices will be fielded; some will be tested as HIPPI attached network storage. Plans are being developed to evaluate shared file systems among multiple client computers with these devices.

#### **5.0 Future Environment**

One of our principal goals for the future is to match massively parallel processing with network attached storage. We have been active with the IEEE Mass Storage Reference Model Committee and other forums and will continue to participate. We intend to deploy Mass Storage technology to selected field sites in order to reduce our network bandwidth requirements. We intend to field an architecture that employs network storage devices which are readily accessible by any of our processors, yet has directly attached storage in those areas where security and access requirements dictate. We will closely monitor research work in these areas as well as that done in tape striping. We will also continue to closely monitor optical tape.

One of our principal lessons learned over the years is that we can no longer afford to enter into a custom development effort for Mass Storage. We must and will rely upon published standards, COTS products, and open systems architectures. We believe that the computer OEMs must accommodate storage product support as a price to do business. Finally, we believe that system reliability, manufacturer warranty, and support costs are just as important as any other acquisition consideration.